

APPLICATION
FOR
UNITED STATES LETTERS PATENT

TITLE: KEYED AUTHENTICATION ROLLOVER FOR ROUTERS

APPLICANT: JACEK SZYSZKO

CERTIFICATE OF MAILING BY EXPRESS MAIL

Express Mail Label No EE 956 370 545 US

I hereby certify that this correspondence is being deposited with the United States Postal Service as Express Mail Post Office to Addressee with sufficient postage on the date indicated below and is addressed to the Commissioner for Patents, Washington, D C 20231

December 11, 2000
Date of Deposit


Signature

Derek W. Norwood
Typed or Printed Name of Person Signing Certificate

KEYED AUTHENTICATION ROLLOVER FOR ROUTERS

BACKGROUND

The invention relates to keyed authentication rollover
5 for routers.

Large networks such as the Internet can be organized
into smaller networks connected by special purpose gateways
known as routers. Hosts and routers, for example, are
presented with Internet Protocol (IP) datagrams addressed
10 to a particular host. Routing is a technique by which the
host or router decides where to send the datagram.

Various routing protocols are available to supply the
information required to perform the routing. For example,
Routing Information Protocol (RIP) routers can exchange
15 topology information with one another. The topology
information defines ways to traverse through networks.
Other devices, such as servers and workstations, may be
connected to the network.

In general, it is important to reduce the likelihood
20 that false protocol messages will be received and processed
by the routers. Routers can use various techniques to
protect themselves against such attacks. Exemplary
algorithms include Message Digest version 4 (MD4) or
version 5 (MD5) algorithms which use encryption-specific

one-way hash functions. According to the MD5 algorithm,
 for example, the routers store a secret key that is used to
 calculate a message digest of the routing information
 placed in each packet. Further details of the MD4 and MD5
 5 algorithms are described in (1) R. Rivest, "The MD5
 Message-Digest Algorithm," MIT Laboratory for Computer
 Science and RSA Data Security, Inc., Network Working Group,
 Request for Comments, RFC 1321 (April 1992) and (2) R.
 Rivest, "The MD4 Message-Digest Algorithm," MIT Laboratory
 10 for Computer Science and RSA Data Security, Inc., Network
 Working Group Request for Comments, RFC 1320 (April 1992).

To increase security, it is desirable to change the
 keys periodically. However, it is important that the
 routers pass information without interruption even if
 15 neighboring routers are not simultaneously configured with
 the new key.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 shows a network.

20 FIGS. 2A and 2B illustrate exemplary encryption key
 lifetimes.

FIG. 3 illustrates exemplary information that is
 stored in memory associated with a router.

FIG. 4 illustrates an exemplary transmission of routing messages according to the invention.

FIG. 5 is a flow chart of a method according to the invention.

5 FIG. 6 illustrates an exemplary format of a routing message.

FIGS. 7A and 7B illustrate routing messages according to an exemplary scenario.

10 DETAILED DESCRIPTION

As illustrated in FIG. 1, an exemplary computer network 10 such as the Internet can include multiple smaller networks connected by routers 12. Smaller networks can include, for example, Ethernet networks 14 and Token-ring networks 16. Networks 18 in a particular geographic area can be connected into a large regional network 20. Other routers (not shown) can pass the information between networks within that area.

Each router 12 has one or more interfaces establishing
20 connections to other routers or networks. Packets are received at input ports and are transmitted from output ports associated with the interfaces. A router 12 examines a received packet of data traveling across the Internet to determine the packet's destination, and the packet is

routed from one router 12 to the next until the packet reaches its destination. Each router maintains a routing table that indicates how to send packets to various destinations. A processor in each router 12 can execute
5 the algorithm discussed below.

Each router 12 can send an advertisement to neighboring routers to inform the neighboring routers of its current routing information. The advertisement can be broadcast or multicast to the neighboring routers and can
10 include one or more routing messages, each of which includes network addresses, cost matrix information or other routing information. The advertisements can be sent on a periodic or other basis. Prior to sending a particular routing message, the router calculates a digest
15 of the routing information using a secret authentication key. The message digest then is transmitted as part of the routing message. The receiving router also uses the authentication key to calculate a digest based on the received routing message and compares its digest to the
20 received digest to authenticate the validity of the received routing message.

Each routing message is assigned a sequence identifier, such as a number, that also is transmitted as part of the routing message. Following receipt of an

initial routing message, a particular router will accept a subsequent routing message only if the sequence number of the later message is higher than the sequence number of the previous routing message.

5 In one implementation, the MD5 algorithm is used to encrypt the routing information, although other algorithms can be used alternatively. Each authentication key is assigned a unique identification and a lifetime, in other words, a time interval during which the key is generally
10 considered to be valid. The authentication key identification can include, for example, a number or other character string. Each key should become valid at a time that differs from the time that any other key on the particular interface becomes valid so that the keys can be
15 sorted by their respective ages.

 The authentication keys and the corresponding lifetimes can be established, for example, by a network administrator. The routers 12 store the keys locally. The keys periodically can be changed by the administrator. For
20 example, the administrator may change the encryption keys once a week, once a month, or according to some other schedule. To limit the amount of administrative overhead required, each router interface can be configured to manage more than one key. Network Time Protocol (NTP) can be used

to synchronize the routers' internal clocks so that, ideally, all neighboring routers 12 begin using a new key at the same time. The keys used on different interfaces for a particular router 12 can be identical or may differ.

5 FIGS. 2A and 2B illustrate alternative ways for specifying the lifetimes for the authentication keys. As shown in FIG. 2A, three keys (Key1, Key2, Key3) are valid during respective time intervals. For example, Key1 is valid during the interval from time t_1 until the time t_3 .
 10 Similarly, Key2 is valid during the interval from time t_2 until the time t_5 . Key 3 is valid during the interval from time t_4 until the time t_6 . In such a situation, there are overlapping periods (indicated by the hatched areas in FIG. 2A) in which a new key and the previous key are both valid.
 15 Furthermore, a single time interval defines the validity of the key at a transmitting router as well as at a receiving router.

FIG. 2B also shows time intervals during which three keys (Key1, Key2, Key3) are valid. In this scenario,
 20 however, different intervals are used to indicate the times when a particular key is valid for use with transmitting routing information and for use with received routing information. For example, Key1 is valid for sending routing information from time t_8 until time t_{10} , whereas

that key is valid for accepting routing information from time t_7 until t_{11} . In the implementation shown in FIG. 2B, the accept interval (hatched area) for Key1 begins before the send interval (non-hatched area) for that key and

5 extends beyond the end of the send interval. Key2 is valid for sending routing information from time t_{10} until time t_{13} , whereas that key is valid for accepting routing information from time t_9 until t_{14} . Similarly, Key3 is valid for sending routing information from time t_{13} until time t_{15} ,

10 whereas that key is valid for accepting routing information from time t_{12} until t_{16} . Thus, as shown in FIG. 2B, the send time interval for a particular key begins substantially at the same time that the send time for the previous key ends. On the other hand, the accept times for sequential keys
15 partially overlap, such that the end of the accept time for a particular key overlaps with the beginning of the accept time for the next key.

As shown in FIG. 3, each router 12 maintains a table of neighboring routers for each of its interfaces. Each
20 neighboring router is identified, for example, by its Internet Protocol (IP) address. Each router also stores the key identification for the last message accepted from each neighboring router 12. A timestamp indicates the most recent time a message received from the neighboring router

was authenticated. Additionally, each router 12 maintains a record of the sequence number identifying the most recently received routing message from each neighboring router. When an authentication key rollover occurs, each
5 router 12 stores the new key identifier in its database.

When a particular router, such as the router 12A (FIG. 4), prepares to send an advertisement about its routing information over a particular interface, it executes the algorithm illustrated by the flow chart of FIG. 5. The
10 interfaces on a particular router 12 can send advertisements independently of one another and need not be synchronized. Initially, the router 12A obtains 100 the current authentication key for the router's interface from its database. The router 12A then determines 102 whether
15 the current key differs from the key used during the previous advertisement for that interface. If the keys differ, then an authentication key rollover has occurred since the last advertisement. A software variable ("old_key") is set 104 to the value of the key used during
20 the previous advertisement ("last_key"). If the determination in block 102 indicates that the keys are the same, then an authentication key rollover has not occurred since the last advertisement. In either situation, the algorithm continues with block 106 in which the router 12A

determines whether all the neighboring routers are configured to use the current key. The determination can be made by reviewing the information stored in the router's table (FIG. 3) and checking the key identifier used in the most recent message accepted from each neighboring router.

If the router 12A determines that all the neighboring routers are configured to use the current key, then the router prepares 108 a particular segment of the routing data and a digest of the segment of the routing data using the current key. The router 12A then sends 110 a message 40 (FIG. 6) that includes a header 42, the particular segment of the routing data 44, the digest 46 of the routing data, the authentication key identification 48 and a sequence number 50. The value of the sequence number for the previous routing message sent by the router 12A over the particular interface is indicated by a variable "SEQ." Thus, the value of the sequence number for the current routing message is set to "SEQ + 1." The value of the variable SEQ then is incremented 112 by one. The cycle of blocks 108, 110 and 112 is continued until all the routing data for the current advertisement has been sent by the router 12A.

After routing messages corresponding to all the routing data have sent to the neighboring routers, a

variable "last_key" that identifies the key used during the previous advertisement is set 114 to the current key.

If (in block 106) the router 12A determines that one or more neighboring routers still are using the old key, then the router 12A prepares 116 a particular segment of the routing data and digests of the segment of the routing data. One digest is calculated using the current key, whereas a second digest is calculated using the old key.

The router 12A sends 118 a first routing message 52 (FIG. 7A) with a format similar to the format of the message 40 shown in FIG. 6. In the first message 52, the digest 46 of the routing data is calculated using the current key and the sequence number 50 for the message is set to "SEQ + 2." The router 12A then sends 120 a second routing message 54 (FIG. 7B). In the second message 54, the digest of the routing data is calculated using the old key and the sequence number 50 is set to "SEQ + 1." The value of the variable SEQ then is incremented 122 by two. The cycle of blocks 116, 118, 120 and 122 is continued until all the routing data for the current advertisement has been sent by the router 12A. After routing messages corresponding to all the routing data have sent to the neighboring routers, the variable "last_key" that

identifies the key used during the previous advertisement is set 114 to the current key.

As indicated by the foregoing discussion, the first routing message 52 prepared using the current key is identified with a sequence number that is higher than the sequence number used to identify the routing message 54 prepared with the old key. FIG. 4 illustrates an exemplary scenario assuming that the previous key is Key1, that the current key is Key2, and that the last routing message sent by the router was identified by the sequence number "N." In that case, the first routing message (prepared with Key2) would be identified by the sequence number "N+2," and the second routing message (prepared with Key1) would be identified by the sequence number "N+1." By sending the first message 52 prior to the second message 54, the amount of processing overhead that must be performed by the receiving routers can be reduced.

In the discussion that follows, it is assumed, for purposes of illustration, that the router 12B (FIG. 4) has successfully performed the new key rollover, but that the router 12C has not yet performed the new key rollover. In such a situation, the transmitting router 12A sends each routing messages twice - first using the new key (Key2) and then using the old key (Key1). The router 12A is

configured to be capable of transmitting routing messages authenticated with the old key even though the normal transmission lifetime for the old key, as indicated by FIG. 2A or 2B, may have expired.

5 A receiving router 12B, 12C will ignore a message if the sequence number associated with the message is not greater than the sequence number of the most recent message processed by that particular router. Furthermore, a receiving router is unable to process a routing message if
10 the identification of the authentication key for the message differs from the authentication key expected by the receiving router.

 Using the example illustrated in FIG. 4, when the router 12C receives a routing message with the new key
15 (Key2), it cannot process the message because the authentication key identification 48 differs from the identification of the key that router 12C expects. That router, however, can accept and process the second message because its sequence number ($N + 1$) is greater than the
20 sequence number (N) of the previously processed message and because the authentication key identification corresponds to the expected key. In contrast, the router 12B accepts the first message with the new key (Key2) because its sequence number ($N + 2$) is greater than the sequence number

(N) of the previously processed message and because the authentication key identification corresponds to the expected key. That router, however, will not process the routing information in the second message because the sequence number (N + 1) of the second message is less than the sequence number (N + 2) of the routing message most recently accepted and processed by that router. Therefore, the router 12B can avoid processing the second message, thereby reducing the total processing overhead.

10 The foregoing techniques can help alleviate problems that may arise when the routers' internal clocks are not perfectly synchronized and/or network management errors occur.

In some situations, it may be desirable for a transmitting router to transmit routing messages using only the new authentication key, even though the transmitting router determines that some of the neighboring routers are not yet using the new key. After the router is powered up, for example, the router will not contain the old key.

20 Therefore, routing messages in the first advertisement after the router is powered up can be sent using only the new key.

Similarly, routing messages in the first advertisement after occurrence of an authentication key rollover can be

sent using only the new key. One rationale for sending routing messages with only the new key in that case can be understood by considering a situation in which there is substantially perfect synchronization among the routers' internal clocks and all neighboring routers rollover to a new key at the same time. In that case, when a particular router is preparing to transmit its next advertisement, it will be unaware that the neighboring routers also have been configured successfully to use the new key. Sending each routing message twice - once with the new key and then with the old key - would be unnecessary. Therefore, in some implementations, the routers are configured not to execute the cycle of blocks 116, 118, 120 and 122 with respect to routing messages that are transmitted as part of the first advertisement following a successful authentication key rollover.

The foregoing techniques can be particularly advantageous when used, for example, with RIP routers described in C. Hedrick, "Routing Information Protocol," STD 34, RFC 1058, Rutgers University (June 1988). However, the techniques can be used with other routers as well.

Various features of the system can be implemented in hardware, software, or a combination of hardware and software. For example, some aspects of the system can be

implemented in computer programs executing on programmable computers. Each program can be implemented in a high level procedural or object-oriented programming language to communicate with a computer system. Furthermore, each such computer program can be stored on a storage medium, such as read-only-memory (ROM), that is readable by a general or special purpose programmable computer, for configuring and operating the computer when the storage medium is read by the computer to perform the functions described above.

Other implementations are within the scope of the following claims.